IN THE CLAIMS:

Please amend claims 23, 31, 39, and 40.

1   1-22. (Cancelled).

1   23.   (Currently Amended) A storage controller, comprising:

2   a destination to store a snapshot from a source;

3   a first write request to write to first data blocks, the first data blocks already cop-

4   ied in the snapshot;

5   the first write request to be written to the source;

6   a second write request to write to second data blocks, the second data blocks not

7   already copied in the snapshot;

8   the second write request being placed into a first in first out queue; and

9   in response to completing the snapshot, the second write request being written

10   from the first in first out queue to the source.

11   snapshot logic;

12   copy logic; and

13   an internal cache;

14   the controller being operable to communicate with a replication manager to re-

15   ceive a snapshot command issued by the replication manager, the snapshot command

16   specifying a range of data bytes of a source volume;

17   the controller being operable to communicate with the replication manager to re-

18   ceive a copy specifying the source volume and a target volume;

19   the controller being operable to receive a write command specifying the source

20   volume;

21   the snapshot logic being operable, in response to the snapshot command, to take a

22   snapshot of the range, the snapshot including a snapshot map and snapshot data, the snap-

23   shot map being stored in a snapshot volume; and

24   ~~the copy logic being operable in response to receiving the copy command to gen-~~
25   ~~erate and send one or more storage device commands to one or more storage devices for~~
26   ~~the source and target volumes to copy data from the source volume directly to the target~~
27   ~~volume without having a file server in the data path, the copy logic using the snapshot~~
28   ~~map and the snapshot data to maintain coherency of the copied data.~~

1   24.   (Previously Presented) The storage device controller of claim 23, wherein the
2   storage device controller is a RAID controller.

1   25.   (Previously Presented) The storage device controller of claim 23, wherein:
2   the range of the source specified by the snapshot command is a first range, and the
3   write command specifies a second range of data bytes of the source volume; and
4   the controller is operable, in response to receiving the write command while the
5   source volume is being copied to the target volume, to hold the write command in the
6   cache, check if the first range overlaps with the second range and, if so, copy the second
7   range from the source volume to the snapshot volume, update the snapshot map, and then
8   allow the write command to write the source volume.

1   26.   (Previously Presented) The storage device controller of claim 23, wherein the rep-
2   lication manager is executed on a file server.

1   27.   (Previously Presented) The storage device controller of claim 26, wherein the file
2   server is connected to a storage area network switch and the file server communicates
3   with the storage device controller through the storage area network switch.

1   28.   (Previously Presented) The storage device controller claim 23, wherein the repli-
2   cation manager is operable to control multiple storage device controllers.

1   29.   (Previously Presented) The storage device controller of claim 23, wherein:

2   the one or more storage device commands include SCSI commands.

1   30.   (Previously Presented) The storage device controller of claim 23, wherein:

2   the controller is operable to send the one or more storage device commands by

3   using one of an in-band protocol or an out-of-band protocol.

1   31.   (Currently Amended) A method comprising:

2   ~~receiving at a storage device controller a snapshot command issued by replication~~

3   ~~manager, the snapshot command specifying a rang of data bytes of a source volume;~~

4   starting a snapshot command from a source to a destination, the snapshot com-

5   mand specifying a range of data bytes of the source;

6   in response to receiving the snapshot command, ~~the~~ a storage device controller

7   taking a snapshot of the range specified using a device control command to control one or

8   more devices on which the source is stored, the snapshot including a snapshot map and

9   snapshot data, and storing the snapshot map and the snapshot data in a cache internal to

10   the storage device controller and snapshot volume, respectively;

11   ~~receiving at the storage device controller a copy command from the replication~~

12   ~~manager, the copy command specifying a copy operation from the source volume to a~~

13   ~~target volume and~~

14   receiving a first write request to first data blocks, the first data blocks already cop-

15   ied in the snapshot;

16   ~~in response to receiving the copy command, the storage device controller generat-~~

17   ~~ing and sending storage device commands to one or more storage devices of the source~~

18   ~~and target volumes to copy data directly from the source volume to the target volume, the~~

19   ~~storage device controller also using the snapshot map and snapshot data to maintain co-~~

20   ~~herency of the copied data.~~

21   writing the first write request to the source;

22   receiving a second write request to write to second data blocks, the second data

23   blocks not already copied in the snapshot;

24 placing the second write request into a first in first out queue; and

25 writing, in response to completing the snapshot, the second request from the first

26 in first out queue to the source.

1 32. (Previously Presented) The method of claim 31, wherein:

2 the storage device is a RAID controller.

1 33. (Previously Presented) The method of claim 31, wherein the range specified by

2 the snapshot command is a first range, the method further comprising:

3 receiving at the storage controller device a write command issued from file sys-

4 tem, the write command specifying a second range of data bytes of the source volume,

5 the write copy command being received while the source volume is being copied to the

6 target volume;

7 in response to receiving the write command, the storage device controller holding

8 the write command in the cache, checking if the first range overlaps with the second

9 range and, if so, copying the second range from the source volume to the snapshot vol-

10 ume, updating the snapshot map, and then allowing the write command to write to the

11 source volume.

1 34. (Previously Presented) The method of claim 31 wherein the replication manager

2 is executed on a file server.

1 35. (Previously Presented) The method of claim 34, wherein the file server is con-

2 nected to a storage area network switch and the file server communicates with the storage

3 device controller through the storage area network switch.

1 36. (Previously Presented) The method of claim 31, wherein the replication manager

2 is operable to control multiple storage device controllers.

1   37.   (Previously Presented) The method of claim 31, wherein the storage device com-

2   mands include a SCSI command.

1   38.   (Previously Presented) The method of claim 31, wherein the storage device com-

2   mands are sent using one of an in-band protocol or an out-of-band protocol.

1   39.   (Currently Amended) A computer-implemented method comprising:

2   starting a snapshot from a source to a destination;

3   receiving a first write request to first data blocks, the first data blocks already cop-

4   ied in the snapshot;

5   writing the first write request to the source;

6   receiving a second write request to write to second data blocks, the second data

7   blocks not already copied in the snapshot;

8   placing the second write request in a first in first out queue;

9   writing, in response to completing the snapshot, the second write request from the

10   first in first out queue to the source.

11   using a replication manager to manage a source storage device controller and a

12   destination storage device controller, the source storage device controller being operable

13   to control access to a source data object and the destination device controller being oper-

14   able to control access to a destination data block, the storage device controllers being op-

15   erable to issue storage device commands;

16   internally generating within the source storage device controller in communica-

17   tion with the replication manager, a snapshot version for each block of the source data

18   object changed by one or more write operations to the block during the course of a copy

19   operations; and

20   copying each block of the source data object to a corresponding block in the des-

21   tination data object in the absence of the snapshot version of the block and otherwise

22   copying the snapshot version of the source data object to block the corresponding block

23   in the destination data object, wherein data is directly transferred between the source and

24 ~~destination storage device controllers without traversing a server operable to process file~~

25 ~~system requests, and wherein coherency of the data transferred between the source and~~

26 ~~destination storage device controllers is maintained without requiring any file system to~~

27 ~~maintain a snapshot map.~~

1 40. (Currently Amended) A system comprising:

2 a destination to store a snapshot from a source;

3 a first write request to write to first data blocks, the first data blocks already cop-

4 ied in the snapshot;

5 the first write request to be written to the source;

6 a second write request to write to second data blocks, the second data blocks not

7 already copied in snapshot;

8 the second write request being placed into a first in first out queue; and

9 in response to completing the snapshot, writing the second write request from the

10 first in first out queue to the source.

11 ~~a replication manager that is operable to issue a snapshot; and~~

12 ~~a storage device controller that is operable (i) to communicate with the replication~~

13 ~~manager to receive the snapshot command and (ii) to receive a copy command, the snap-~~

14 ~~shot command specifying a range data bytes of a source volume, the copy command~~

15 ~~specifying the source volume and a target volume; wherein~~

16 ~~the controller is operable to receive a write command specifying the source vol-~~

17 ~~ume;~~

18 ~~the controller is operable, in response to receiving the snapshot command, to take~~

19 ~~a snapshot of the range, the snapshot including a snapshot map and a snapshot data, the~~

20 ~~snapshot map being stored in a cache internal to the storage device controller and the~~

21 ~~snapshot data being stored to a snapshot volume; and~~

22 ~~the controller is operable, in response to receiving the copy command, to generate~~

23 ~~and send one or more storage device commands to one or more storage devices for the~~

24 ~~sources and target volumes to copy data from the source volume directly to the target~~

7

25 ~~volume without having a file server in the data path, the controller using the snapshot~~

26 ~~map and the snapshot data to maintain coherency of the copied data.~~

1 41. (Previously Presented) The system of claim 40, wherein the replication manager

2 is executed on a file server and is operable to control the source storage device controller

3 and one or more other storage device controllers.

1 42. (Previously Presented) The system of claim 40, further comprising a list of source

2 data blocks to be copied that are reordered to increase copy speed.

1 43. (Previously Presented) The system of claim 42, wherein the list of blocks to be

2 copied is buffered while the storage device controller awaits further copy commands.

1 44. (Previously Presented) The system of claim 40, wherein the replication manager

2 is operable to inserted control data before and after a source data block being copied.

1 45. (Previously Presented) The system of claim 40, wherein the replication manager

2 is operable to specify a block size so that the storage device controller writes fixed-size

3 blocks.

1 46. (Previously Presented) A method, comprising:

2 receiving a write request, wherein the write request has a range of data bytes to

3 modify;

4 determining the range of bytes is within a snapshot range, wherein a snapshot is a

5 process of copying data to a new location before the data is modified by a write opera-

6 tion;

7 determining the range of bytes has not been snapshotted;

8 in response to determining the range of bytes has not been snapshotted, copying

9 the range of bytes from a source volume to a snapshot volume;

10   updating a volume snapshot map, wherein the snapshot map determines which

11 blocks are located in the snapshot volume, in response to copying the range of bytes;

12   modifying the range of bytes of data from the write request, in response to copy-

13 ing the range of bytes; and

14   copying the range of bytes from the source volume to a target volume using the

15 volume snapshot map and data stored in the snapshot volume, and copying without hav-

16 ing a file server in the path.